
Generalised Entropy MDPs and Minimax Regret

Emmanouil G. Androulakis
Göteborgs Universitet
Gothenburg, Sweden
emmanouil.an@outlook.com

Christos Dimitrakakis
Chalmers University of Technology
Gothenburg, Sweden
chrdimi@chalmers.se

Abstract

Bayesian methods suffer from the problem of how to specify prior beliefs. One interesting idea is to consider worst-case priors. This requires solving a stochastic zero-sum game. In this paper, we extend well-known results from bandit theory in order to discover minimax-Bayes policies and discuss when they are practical.

1 Introduction

In this work, we consider the problem of a Bayesian agent interacting with a Markov decision process (MDP). However the agent is unsure of how to select its prior distribution and so it prefers a choice that is safe against a potentially adversarial Nature. The problem is how to select such a policy in a computationally efficient manner. We first recall the definition of an MDP.

Definition 1. *A Markov decision process $\mu \in \mathcal{M}$ on a state-action space $\mathcal{S} \times \mathcal{A}$ is a tuple $\langle \mathcal{S}, \mathcal{A}, P, \rho, T \rangle$ where \mathcal{S} is a set of states, \mathcal{A} is a set of actions, P is a transition kernel, such that*

$$s_{t+1} \mid s_t = s, a_t = a \sim P^{s,a},$$

where $r_t = \rho(s_t)$ and T is a (potentially random) horizon.

The agent's utility is an additive function of individual rewards $r_t \in \mathbb{R}$

$$U \triangleq \sum_{t=1}^T r_t. \quad (1)$$

For simplicity, we can assume that the reward function is known to the agent and then for a finite state space ρ can be taken to be a fixed vector. For any MDP $\mu \in \mathcal{M}$ and policy $\pi \in \Pi$, the expected utility is $\mathbb{E}_\mu^\pi(U)$, while the conditional expected utility is called the *value function*:

$$V_\mu^\pi(s) \triangleq \mathbb{E}_\mu^\pi \left(\sum_{t=1}^T r_t \mid s_t = s \right). \quad (2)$$

For finite MDPs, and a $(1 - \gamma)$ -geometrically distributed horizon T , the value function can be written as a vector $V_\mu^\pi = (I - \gamma P_\mu^\pi)^{-1} \rho$, where P_μ^π is the Markov chain induced on the MDP by the policy π . Since we are uncertain about μ , we can instead define a prior distribution ξ on \mathcal{M} . Then

$$V_\xi^\pi \triangleq \int_{\mathcal{M}} V_\mu^\pi d\xi(\mu), \quad (3)$$

will denote the value function under the particular distribution ξ on the MDPs. Finally, for a given probability measure β on \mathcal{S} , which can be taken to represent a starting state distribution, we define the utility of a particular policy π to be:

$$U(\xi, \pi) \triangleq \mathbb{E}_\xi^\pi(\beta^\top V_\mu^\pi) = \int_{\mathcal{M}} \beta^\top V_\mu^\pi d\xi(\mu). \quad (4)$$

There are two possible ways to interpret the measure ξ , depending on how it is chosen. If ξ is selected by the agent selecting π , then it corresponds to the subjective belief of the decision maker about which is the most likely MDP *a priori*. Then $U(\xi, \pi)$ corresponds to the expected utility of a particular policy under this belief. Let

$$U^*(\xi) \triangleq \sup_{\pi \in \Pi} U(\xi, \pi) \quad (5)$$

denote the Bayes-optimal utility for a belief. We recall the fact that this is a convex function [c.f. DeGroot, 1970]. By definition, and due to convexity, the following bounds hold:

$$U(\xi, \pi) \leq U^*(\xi) \leq \int_{\xi} U^*(\mu) d\xi(\mu), \quad \forall \pi \in \Pi. \quad (6)$$

In the above, the left hand side is the utility of an arbitrary policy, while the right side can be seen as the expected utility we would obtain if the true MDP was revealed to us.

The second view of ξ is to assume that the MDP is *actually* drawn randomly from the distribution ξ . If this is known, then the subjective value of a policy is equal to its true expected value. However, it is more interesting to consider the case where Nature selects some ξ in an arbitrary way from a *set* of possible priors Ξ . Then we wish to find a policy π^* achieving:

$$\max_{\pi \in \Pi} \min_{\xi \in \Xi} U(\xi, \pi). \quad (7)$$

One basic open question is whether the maximum exists. This is answered in the affirmative when the game between nature and the agent has a value, i.e.

$$U^* = \sup_{\pi \in \Pi} \inf_{\xi \in \Xi} U(\xi, \pi) = \inf_{\xi \in \Xi} \sup_{\pi \in \Pi} U(\xi, \pi) = U_*. \quad (8)$$

Let π^* and ξ^* be the maximin policy and minimax prior respectively. If the game has a value then there exists an equalising policy which is optimal for some belief ξ^* , and vice versa. A sufficient condition for this to occur is for U^* to be convex and differentiable everywhere [c.f. Grünwald and Dawid, 2004]. In order to study when this can occur, we first go over a couple of well-known facts.

2 Existence of maximin policies

Definition 2 (Policy). Let \mathcal{H} be the set of all histories (s_t, a_{t-1}) . A (stochastic) policy π is a set of probability measures $\{\pi(\cdot | h) \mid h \in \mathcal{H}\}$ on the set of actions \mathcal{A} . We denote the set of all (history-dependent, stochastic) policies by Π^S .

Definition 3 (Deterministic policy). A policy is deterministic if, for each sequence s^t, a^{t-1} , there exists an action $a \in \mathcal{A}$ such that $\pi(a_t = a \mid s_t, a_{t-1}) = 1$. We denote the set of deterministic policies by Π^D .

Definition 4 (Memoryless policy). A policy is memoryless (or reactive) if, for all sequences s_t, a_{t-1} , we have $\pi(a_t = a \mid s_t, a_{t-1}) = \pi(a_t = a \mid s_t)$. We denote the set of memoryless (stochastic) policies by Π_1^S .

The set of memoryless deterministic policies is denoted by Π_1^D . Obviously, $\Pi_1^D \subset \Pi^D \subset \Pi^S$ and $\Pi_1^D \subset \Pi_1^S \subset \Pi^S$.

Definition 5 (Mixed policy). A mixed policy is a probability measure over policies. If Π is a set of base policies, we denote the set of probability measures over Π by $\Delta(\Pi)$.

Fact 1. For any MDP μ there exists a deterministic, memoryless policy that is optimal, i.e. $U^*(\mu) = \sup_{\pi \in \Pi} U(\mu, \pi) = \max_{\pi \in \Pi_1^D} U(\xi, \pi)$.

Fact 2. For any distribution ξ over MDPs, there exists a deterministic, history-dependent policy that is optimal, i.e. $U^*(\xi) = \sup_{\pi \in \Pi} U(\xi, \pi) = \max_{\pi \in \Pi^D} U(\xi, \pi)$

A well-known game theoretic result is that an equalising policy can always be found in $\Delta(\Pi)$ when Π is finite. However for fixed, finite T , there may not exist a deterministic equalising policy. Then the number of possible policies is finite and consequently U^* is piecewise-linear.

Remark 1. If Π^D is finite, there exists a policy $\pi^* \in \Pi^S$ achieving the value of the game.

Proof. For any mixed policy $\delta \in \Delta(\Pi^D)$, there exists an equivalent stochastic policy in Π^S . This can be constructed by augmenting the state space to include the outcomes of fair coins. Now note that there exists an optimal mixed strategy $\delta^* \in \Delta(\Pi^D)$ achieving the value of the game, as the number of finite-horizon policies is finite. But there also exists a stochastic policy with the same distribution for all histories $h \in \mathcal{H}$. ■

Grünwald and Dawid [2004] make some interesting connections between maximum entropy and robust Bayesian decisions. In particular, they define the generalised entropy of a distribution ξ to be minimum loss $H(\xi) \triangleq \inf_{\pi} L(P, \pi)$ achievable in a game between nature and a decision maker. In our setting, it is natural to consider the following two loss functions:

$$L_1(\xi, \pi) = \beta^\top (V_\xi^* - V_\xi^\pi), \quad L_2(\xi, \pi) = \beta^\top \mathbb{E}_\xi(V_\mu^* - V_\mu^\pi), \quad (9)$$

where $\beta \in \mathbb{R}^{|\mathcal{S}|}$ is a distribution on the states. These corresponds to the regret of π relative to the ξ -optimal policy and to the oracle policy respectively. Now, let Ξ be a set of probability distributions on \mathcal{M} . One idea is to try and guard against the worst-case prior in a restricted set Ξ_ϕ , by constraining the expectation of a statistic $\Phi : \mathcal{M} \rightarrow \mathbb{R}^k$ under the prior to be equal to the observed value of the statistic, ϕ .

$$\Xi_\phi = \{ \xi \in \Xi \mid \mathbb{E}_\xi(\Phi) = \phi \}. \quad (10)$$

One possibility is to use the cumulative state distribution for a particular policy π , i.e.

$$\Phi(\mu) = (I - \gamma P_\mu^\pi)^{-1}, \quad (11)$$

where a common choice for π is the optimal policy for MDP μ , $\pi^*(\mu)$, used for example in Mannor and Shimkin [2003]. In that case it is easy to see that $V_\xi^* = \mathbb{E}_\xi(\Phi)\rho$. A sufficient condition for a policy π to be robust Bayes against Ξ_ϕ is for it to be linear, that is $L(\mu, \pi) = \alpha_0 + \alpha^\top \Phi(\mu)$ for all $\mu \in \mathcal{M}$. Then it is also true [see Grünwald and Dawid, 2004, Theorem 7.1] that π is an equalising policy against Ξ_ϕ .

How can Ξ_ϕ be calculated? When the set of MDPs is finite, then it is defined through the linear equation $\Xi_\phi = \left\{ \xi \mid \sum_{\mu} \xi(\mu) \Phi(\mu) = \phi \right\}$. Given, then, a sequence of observations, a statistic, and a resulting set of priors Ξ_ϕ , one important question is how we can efficiently calculate such policies. We explain this in the following section.

3 Calculating robust policies

One potential solution involves finding the minimax prior ξ^* and then the policy that is Bayes-optimal with respect to it. In previous work Koolen [2006] has shown that finding ξ^* can be found via a concave-linear optimisation, for the ‘truth-finding’ game. However, this is not generally true. However, it has been shown by Freund and Schapire [1999], that multiplicative weights algorithm can be used for zero-sum games, as long as an oracle that can compute best responses is available. In our setting, this would correspond to nature having the ability to efficiently construct a worst-case MDP given a policy. On the other hand Vrieze and Tijs [1982] show that the value in stochastic zero sum games can be achieved asymptotically via fictitious play even when the game matrix is known approximately, as long as the approximation converges to the true game matrix.

One idea is to apply results from the experts literature, such as the weighted majority algorithm (WMA). We start by assuming that in each round k the Decision Maker has full

access to the information regarding the rewards of the past round. That means that she can observe the outcomes of *all* the policies that were available previously.

Assume that $|\mathcal{M}| = M$ is finite¹. Then let \mathbf{v}_π be the $1 \times M$ vector of values for policy π for the given set of Markov decision problems \mathcal{M} :

$$\mathbf{u}^\pi = (U(\mu, \pi))_{\mu \in \mathcal{M}}.$$

Assume that $|\Pi| = N$ is finite and denote by \mathbf{u}_ξ the $1 \times N$ vector of values for each policy

$$\mathbf{u}_\xi = (U(\xi, \pi))_{\pi \in \Pi}.$$

To apply WMA (Alg. 2) we execute policies in rounds. At each round, we select a policy $\pi_{(k)}$ from a distribution q_k , and nature chooses some prior $\xi_{(k)}$. Moreover, denote by $x_{\pi_i, k}$ the *total realized reward* obtained by following policy π_i , $i = 1, 2, \dots, N$, in the k -th round. Each $x_{\pi_i, k}$ is a random variable that has an expected value, equal to $U(\xi, \pi_i)$. Finally, denote by $\mathbf{x}_{(k)}$ the vector of sampled utilities of all policies up to round k :

$$\mathbf{x}_{(k)} = (x_{\pi_1, k}, x_{\pi_2, k}, \dots, x_{\pi_N, k}).$$

Algorithm 1 WMA

Input: A set of policies π , with $|\pi| = N$; a set of weights $w_{(k)} = (w_{i,k})_{i=1}^N$; a learning rate $0 < \ell \leq 1/2$. **Initialize:** $w_{i,1} = 1$.

For each round k :

- 1: DM(Decision Maker) normalizes the weights to get a distribution $q_k = \frac{w_{(k)}}{\sum_{i=1}^N w_{i,k}}$
- 2: DM selects $\pi_{(k)}$ among π_i , $i = 1, 2, \dots, N$ according to the distribution q_k
- 3: Nature chooses $\xi_{(k)} \in \operatorname{argmin}_{\xi_{(k)}} \mathbb{E}_{q_k} [\mathbf{u}_{\xi_{(k)}}]$
- 4: DM receives reward $x_{k, \pi_{(k)}}$ and calculates $U(\xi_{(k)}, \pi_i)$ for all policies $\pi_i \in \pi$
- 5: DM calculates the next set of weights for $i = 1, \dots, N$:

$$w_{i,k+1} = (1 + \ell U(\xi_{(k)}, \pi_i)) w_{i,(k)}$$

3.1 Analysis

The main issue is the computation of the value $U(\xi_{(k)}, \pi_i)$, which is used in steps 3 and 4. Depending on the sizes of the policy and MDP space accurate or approximate values for the quantity $U(\xi_{(k)}, \pi_i)$ can be obtained.

When N (number of policies) and M (number of MDPs) are small. Then we can retrieve the expected value of each policy $V_\mu^{\pi_i} = (I - \gamma P_\mu^{\pi_i})^{-1} \rho$, where $P_\mu^{\pi_i}$ is the kernel. The inverse operator will require $O(|\mathcal{S}|^3)$. The expected reward for sampling a policy π from the distribution q_k is

$$\mathbb{E}_{\pi \sim q_k} [x_{k, \pi}] = \mathbf{x}_{(k)} \cdot q_k.$$

The total expected reward over all rounds is therefore

$$\mathcal{V}_{\text{WMA}}^{(K)} \triangleq \sum_{k=1}^K \mathbf{x}_{(k)} \cdot q_k.$$

Theorem 1 (Arora et al. [2012]). *The Multiplicative Weights algorithm guarantees that after K rounds, for any distribution \mathcal{Q} on the decisions, it holds:*

$$\mathcal{V}_{\text{WMA}}^{(K)} \geq \sum_{k=1}^K (\mathbf{u}_{\xi_k} - \ell |\mathbf{u}_{\xi_k}|) \cdot \mathcal{Q} - \frac{\log_e N}{\ell}$$

where $|\mathbf{u}_{\xi_k}|$ is the vector obtained by taking the coordinate-wise absolute value of the vector containing the expected values $V_{\xi_k}^{\pi_i}$.

¹If $M = \infty$, a grid over the MDPs can be used to obtain a finite set. Then the approximation error for an MDP that has an ϵ -close transition matrix and mean reward from an MDP on the grid, will be bounded by $\epsilon/(1 - \gamma)^2$.

If M is small, but N is large then computation of $U(\xi_{(k)}, \pi_i)$ becomes difficult. We can approximate the true value of the expected values, though, by using a Monte Carlo sampling of S iterations. We substitute $U(\xi_{(k)}, \pi_i)$ with its estimator $\hat{U}(\xi_{(k)}, \pi_i, S)$ (lines 3,4 & 5 of algorithm WMA) and we call this modification of the algorithm as WMA-SR. Moreover, for each round k , we need to introduce an estimation error term $\mathcal{E}_{(k)}^{(S)}$ (which depends on the number S of Monte Carlo iterations), since now the weights are updated by using approximations and not the true values.

Algorithm 2 WMA-SR

Input: A set of policies π , with $|\pi| = N$; a set of weights $w_{(k)} = (w_{i,k})_{i=1}^N$; a learning rate $0 < \ell \leq 1/2$. **Initialize:** $w_{i,1} = 1$.

For each round k :

- 1: DM(Decision Maker) normalizes the weights to get a distribution $q_k = \frac{w_{(k)}}{\sum_{i=1}^N w_{i,k}}$
- 2: DM selects $\pi_{(k)}$ among π_i , $i = 1, 2, \dots, N$ according to the distribution q_k
- 3: Nature chooses $\xi_{(k)} \in \operatorname{argmin}_{\xi_{(k)}} \mathbb{E}_{q_k} [\mathbf{u}_{\xi_{(k)}}]$
- 4: DM receives reward $x_{k,\pi_{(k)}}$ and calculates $\hat{U}(\xi_{(k)}, \pi_i, S)$ for all policies $\pi_i \in \pi$
- 5: DM calculates the next set of weights for $i = 1, \dots, N$:

$$w_{i,k+1} = \left(1 + \ell \hat{U}(\xi_{(k)}, \pi_i, S)\right) w_{i,(k)}$$

The value earned by using WMA-SR, over all rounds is

$$\mathcal{V}_{\text{WMA-SR}}^{(K)} \triangleq \mathbb{E} \left(\sum_{k=1}^K \mathbf{x}_{(k)} \cdot q_k \right) = \sum_{k=1}^K \mathbf{u}_{\xi_{(k)}} \cdot q_k$$

where $q_k = (q_{k,1}, \dots, q_{k,N})$.

First, we prove a lemma for the approximated expected values.

Lemma 1. *Assume that all policy rewards lie in $[-1, 1]$. Let $0 < \ell \leq \frac{1}{2}$. Then after K rounds, it holds:*

$$\mathcal{V}_{\text{WMA-SR}}^{(K)} \geq \sum_{k=1}^K \hat{U}(\xi_{(k)}, \pi_i, S) - \ell \sum_{k=1}^K |\hat{U}(\xi_{(k)}, \pi_i, S)| - \frac{\log_e N}{\ell}$$

for all $i = 1, 2, \dots, N$.

Observe that if the rewards are not stochastic, then Lemma 1 is reduced to the standard expert setting [Arora et al., 2012, Freund and Schapire, 1999, c.f].

Theorem 2. *Assume that all policy rewards lie in $[-1, 1]$. Let $0 < \ell \leq \frac{1}{2}$. Let $\varepsilon > 0$. Then after K rounds, for the total expected rewards, it holds:*

$$\mathcal{V}_{\text{WMA-SR}}^{(K)} \geq \sum_{k=1}^K U(\xi_{(k)}, \pi_i) - \ell \sum_{k=1}^K |U(\xi_{(k)}, \pi_i)| - \frac{\log_e N}{\ell} - \sum_{k=1}^K \mathcal{E}_{(k)}^{(S)}$$

for all $i = 1, 2, \dots, N$, where $\xi_{(k)} = (\xi_{1,k}, \dots, \xi_{M,k})$, $\sum_m \xi_{m,k} = 1$, $q_k = (q_{1,k}, \dots, q_{N,k})$, $\sum_i q_{i,k} = 1$, $\mathcal{E}_{(k)}^{(S)}$ is the error term of the k -th round (and S denotes the number of Monte Carlo simulations):

$$\mathcal{E}_{(k)}^{(S)} = \left| \mathbf{u}_{\xi_{(k)}} \cdot q_k - U(\xi_{(k)}, \pi_i) + \ell |U(\xi_{(k)}, \pi_i)| + \frac{\log_e N}{\ell} - \left(\hat{\mathbf{u}}_{\xi_{(k)}} \cdot q_k - \hat{U}(\xi_{(k)}, \pi_i, S) + \ell |\hat{U}(\xi_{(k)}, \pi_i, S)| + \frac{\log_e N}{\ell} \right) \right|$$

and

$$\mathbb{P} \left(\left| \sum_{k=1}^K \mathcal{E}_{(k)}^{(S)} \right| < \varepsilon \right) \geq 1 - 2 \exp \left(-k \frac{\varepsilon^2}{2} \right).$$

The error bound is retrieved by applying Azuma's Lemma, since $\hat{U}(\xi_{(k)}, \pi_i, S) - U(\xi_{(k)}, \pi_i)$ is a martingale difference and all rewards lie in $[-1, 1]$. We can also obtain a result for a distribution \mathcal{P} over π_i 's, $i = 1, 2, \dots, N$.

Corollary 1. *After K rounds, for any distribution $\mathcal{P} \in \mathbb{R}^{N \times 1}$ on the decisions, it holds:*

$$\sum_{k=1}^K \mathbf{u}_{\xi_{(k)}} \cdot q_k \geq \sum_{k=1}^K (\mathbf{u}_{\xi_{(k)}} - \ell |\mathbf{u}_{\xi_{(k)}}|) \cdot \mathcal{P} - \frac{\log_e N}{\ell} - \sum_{k=1}^K \mathcal{E}_{(k)}^{(S)}$$

where $|\mathbf{u}_{\xi_{(k)}}|$ is the vector obtained by taking the coordinate-wise absolute value of $\mathbf{u}_{\xi_{(k)}}$.

Definition 6. *The regret of the learning algorithm against the optimal distribution $\mathcal{P}^* \in \arg\max_{\mathcal{P}} \{\mathbf{u}_{\xi_{(k)}} \cdot \mathcal{P}\}$ is*

$$B(K) = \sum_{k=1}^K \mathbf{u}_{\xi_{(k)}} \cdot \mathcal{P}^* - \sum_{k=1}^K \mathbf{u}_{\xi_{(k)}} \cdot q_k$$

Corollary 1 can be used in order to bound the regret. To that end, we first need the following theorem.

Theorem 3. *After K rounds of applying the modified weighted majority algorithm WMA-SR, for any distribution \mathcal{P} it holds:*

$$\sum_{k=1}^K \mathbf{u}_{\xi_{(k)}} \cdot \mathcal{P} - \sum_{k=1}^K \mathbf{u}_{\xi_{(k)}} \cdot q_k \leq 2\sqrt{\log_e NK} + \sum_{k=1}^K \mathcal{E}_{(k)}^{(S)}$$

Corollary 2. *When algorithm WMA-SR is run with parameter $\ell = \sqrt{\frac{\log_e N}{K}}$ then the regret of the algorithm is bound by $B(K) \leq 2\sqrt{\log_e NK} + \sum_{k=1}^K \mathcal{E}_{(k)}^{(S)}$.*

One can also show that the algorithm converges by dividing the time into epochs. A choice of epochs that gives convergence with probability one is to define the length of each epoch as $T_k = k^2$ and then proceed similarly to Freund and Schapire [1999] to show the following

Theorem 4. *Suppose we repeat the game for an unbounded number of rounds. Then for the regret of the algorithm it holds: $\mathbb{P} \left[\frac{B(K)}{K} \leq \epsilon \right] = 1$ for all but a finite number of values of K and for every $\epsilon > 0$.*

View as a bandit problem. To avoid uniformly sampling all policies, we could cast this as a contextual bandit problem, mapping the prior ξ_k selected at each round by nature to the context. With a slight modification of the linear context bandit algorithm presented in Auer [2002], we then obtain a similar bound, which is however linear in the number of policies, making this approach impractical. Since our estimates converge to the values of the policies, we can recover the value of the game if Nature uses fictitious play with respect to our estimates [Vrieze and Tijs, 1982].

4 Conclusion

We have discussed the links between robust reinforcement learning and maximin policies. In particular, an interesting idea to explore is to use maximin policies against a constrained set of priors Ξ_ϕ . Such as set is easy to define for a finite number of MDPs. However, to put the idea into practice we also need to establish computational procedures for approximately calculating such policies. Although this seems to be achievable for small problems, there does not appear to be a useful procedure for larger ones. One potential direction would be to choose statistics Φ that inherently make the problem amenable to simple solutions.

References

Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.

- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2002.
- Morris H. DeGroot. *Optimal Statistical Decisions*. John Wiley & Sons, 1970.
- Yoav Freund and Robert E Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1):79–103, 1999.
- Peter D. Grünwald and A. Philip Dawid. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *Annals of Statistics*, 2004.
- Wouter Michiel Koolen. Discovering the truth by conducting experiments. Master’s thesis, University of Amsterdam, 2006.
- Shie Mannor and Nahum Shimkin. The empirical Bayes envelope and regret minimization in competitive Markov decision processes. *Mathematics of Operations Research*, 28(2):327–345, 2003. ISSN 0364-765X.
- O.J. Vrieze and S.H. Tijs. Fictitious play applied to sequences of games and discounted stochastic games. *International Journal of Game Theory*, 11(2):71–85, 1982. ISSN 0020-7276. doi: 10.1007/BF01769064. URL <http://dx.doi.org/10.1007/BF01769064>.